Gigabits to the Desktop: Installing Tomorrow's Networks Today

Timothy C. Kuhfuss and Paul T. Phillips
Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439-4844
Internet: kuhfuss@anl.gov
Internet: phillips@anl.gov

Abstract

The dream of gigabits to the desktop is becoming reality. Computer scientists and engineers are now building gigabit architectures in the labs. Leading edge computing applications are starving for bandwidth, desktop workstations are quickly evolving, and gigabit network architectures are being developed to support high speed communications. This paper will briefly discuss progress in these areas, but will primarily focus on the foundation of gigabits to the desktop, the physical layer. It will discuss gigabit physical media options and recommendations for selecting and installing appropriate hardware infrastructures. Lastly, we will discuss Argonne's experience in prototyping physical layer hardware.

1 Introduction

Argonne is one of the U. S. Department of Energy's world class research institutions. Leading edge computing tools and networks allow Argonne to maintain and enhance this reputation. One current effort to deploy leading edge tools is the Argonne "Gigabits to the Desktop" project. While delivering and using gigabits to the desktop is little more than a hope at this time, this paper will discuss the hurdles to achieving it and how to tear down as many hurdles as possible.

Under this project, four distinct areas are being investigated and enhanced. This paper will briefly discuss the applications and tools that we see driving the requirement for gigabits to the desktop. It will touch on a functional description of our "ideal" workstations, architectures and the candidates for the next generation network capable of delivering gigabits. Lastly, it will provide an in-depth analysis of physical layer options and attempt to prove that this area, while the least risky, must be done properly, with the proper media.

This paper assumes one important point. It assumes that bandwidth is essentially free. We will discuss network architectures and physical installation recommendations which have a fixed cost. However in a campus environment, there is no marginal cost for additional packets on these networks once the network infrastructure is installed. This point is important when extrapolating our conclusions to the wide area. The marginal cost of a packet sent to a commercial network is usually non zero. This fact may prove to be a great hindrance in migrating the applications mentioned beyond the organizational boundaries.

The overall goal of Argonne computing is to push network technology, computing tools and communication systems to their limits enabl ing Argonne researchers to perform "world class science". Our plan is to create the ideal scientific work environment with the use of distributed supercomputers, fast workstations, seamless high speed networks and efficient scientific applications. These tools can enable our researchers to study classes of problems not possible previously.

Argonne is delivering this work environment by attacking four areas. The top down view of this environment is applications, desktop architectures, networks and physical infrastructure. For the most part the applications are well known, the desktop architectures are rapidly evolving, the network options are narrowing down to a few and the physical layer is stabilizing. To date, this project has put the majority of its effort into the physical infrastructure. Our feeling is that a solid physical base will enable us to attack other areas more easily. We also realize that poor physical installation s will cripple future progress of this project as a whole.

2 Application Concentration

The goal of the application concentration is to provide and demonstrate computing tools to enhance the research done at Argonne. Many applications are now well known; only now possible with existing desktop architectures. Many more have yet to be considered and will be enabled by progress in the development of current tools. The delivery of a low latency, high bandwidth network enables the following applications:

Real Time Visualization - The display of high resolution images at 30 frames per second can generate nearly a gigabit of data per second (1024x12 80 pixels/frame x 24 bits/pixel x 30 frames/second $\cong$ 950 Mbps non-interlaced). Compression techniques can reduce the amount of required network traffic in many cases.

Virtual Reality - The ra pid display of high definition, stereoscopic images required for virtual reality requires over a gigabit per second of data capacity if the images are displayed in real time.

Distributed Virtual Reality - Real time interaction among virtual reality devices can require large data transfer rates if the virtual reality viewers are driven from central servers. Preloading the VR environment and only exchanging coordinate changes is one way around this problem.

Distributed Collaboration - Many researcher s desire the ability to collaborate with their colleagues located throughout the world. Conferencing tools such as video and audio exchange s and distributed white boards are essential. Remote display of scientific images is also critical. The ability to manipulate the image locally while an instructor points out relevant events is highly desirable.

Broadcast Video - Displaying video from central servers will require great amounts of bandwidth at the server end and modest amounts at the desktop. Control signals (RW,Stop,FF) will have to be sent in real time. Preloading the video to the desktop is one way to work around the bandwidth requirement. Unfortunately this work around requires massive amounts of local data storage. Videos can be downloaded to the desktop as fast as possible and manipulated locally.

Massive (Petabyte) Databases - Scientific databases are becoming larger and more critical. Many now download the entire database and do local inquiries. The most efficient way to run these databases is to locate them centrally and allow remote queries. Single database inquiries can generate a large amount of response data to be transferred.

Aggregate Bandwidth (at the desktop) - The least flashy application requiring large amounts of bandwidth is the aggregation of all the applications mentioned above. Even if the individual applications are optimized to reduce network impact, combining a set of the ones mentioned above will require gigabit connections at the desktop.

Tuning these applications using techniques mention ed above can minimize overall network impact. However, all this tuning often trades off local processing with network bandwidth. Much debate of the merits of large pipes (high bandwidth networks) vs. local compression has been going on over past years. Our feeling is that using compression techniques in conjunction with large pipes will optimize current and future application development.

3 Desktop Architectures Concentration

The goal of the desktop architectures is to push workstation manufacturers to create appropriate desktop hardware, optimized for high speed network access, for scientific work.

We have found that as a minimum the desktop machines must have a high performance processor for local computation, a high resolution display for scientific graphics and easy display of disparate application windows, compact disk quality sound, video capture capability for conferencing and input/output of video for transport and archiving, and support high speed networking.

In general, a modular architecture is desirable so that the scientist can grow their systems as they need features and can afford them. Another desirable feature of a desktop architecture is intelligent peripherals. A good deal of processing of peripheral data (network, video, etc.) can be done on the interface cards or peripherals and should be.

Another appropriate topic when discussing desktop architectures is the fact that high speed network connections may not be optimized using current system bus architectures. Current architectures separate the Input/Output portion of the machine from the central processor and memory. With this architecture, reading large amounts of data to memory requires a great deal of CPU and system bus access. This has proven in the past to be the biggest bottleneck for tuning machines to use high speed networks.

Computer scientists are now investigating using the network as a system bus. The latency and bandwidth characteristics of many emerging networks is approaching system bus performance levels. This enables the option of creating desktop machines by plugging appropriate smart peripherals to a desktop network bus. Argonne is currently investigating and implementing this type of architecture to a massively parallel computer but would like to see it migrate to a desktop.

4 Network Layer Concentration

The goal of the network concentration is to investigate and exploit network architectures to enhance ANL applications. Specifically we would like low latency, high bandwidth on demand, multimedia networks to every desktop at the price of a current Ethernet connection.

The simplistic view of bandwidth vs. latency is that bandwidth deals with how much data gets to the desktop; latency deals with how quickly. The majority of the work done to date has concentrated on bandwidth but ignored latency. This is unfortunate because any of the applications mentioned above with a real time component require quick communication as well as high bandwidth.

One often used analogy of automobiles describes typical networks as cars driving coast to coast over the highway. The auto equivalent of a high speed network is not a faster car , as many people believe, but a wider road. The speed that the first car (or a bit) can go is a relatively fixed cost. The amount of people (or data) delivered is a function of the road size (or bandwidth). Buses would be the equivalent of large packet networks; a lot of cars would be the fixed cell equivalent.

The issue of latency can best be described as tolls. A car should go a fixed speed of approximately 55-65 mph. Similarly, bits will travel at approximately near the speed of light over copper or fiber. The car will be slowed down by things like tolls and gas stations. Bits will be slowed down by things like switches and routers.

In the early 1980s there was a bloody desktop war over the "best megabit network". Ethernet and Token Ring were the leading battlers with AppleTalk and Arcnet contributing to the mess. The battle continues to go on with a comfortable lead by Ethernet.

In early 1993, a new war began over the best gigabit architecture. A short list of possible candidates and a quick assessment follows:

100 M Ethernet - Ethernet running at 100 Mbps is current ly a hot topic. The reason for this effort is to offer a high speed LAN at a fraction of the FDDI costs. Unfortunately, fragmented standards efforts are impeding progress of this standard.

ATM - Asynchronous Transfer Mode (ATM) is a cell switched architecture promising variable bandwidth on demand from 51 Mbps to over 2.4 Gbps, quality of service guarantees, virtual network support and many other features. If standardized and implemented correctly , this standard promises to be the LAN and WAN of the foreseeable future.

Fibre Channel - Fibre Channel System (FCS) is a recent channel standard which, like HiPPI, can be switched to create a network. Prominent network speeds are 266 Mbps (quarter speed) and 1 Gbps. In addition to interconnecting machines and peripherals, FCS is being used to cluster high performance workstations for distributed computing.

FDDI, CuDDI, FFOL - Fiber Distributed Data Interface (FDDI) and its copper equivalent are currently the only solutions for more than 20 Mbps. FDDI is a shared 100 Mbps token passing network. FDDI has a version which supports isochronous data such as voice, FDDI-II. FDDI Follow On LAN is a standard for a gigabit version of FDDI. To date, little progress has been made in this standard and may continue to progress slowly unless all of the other network architectures fail.

HiPPI - High Performance Parallel Interface (HiPPI) is a copper based channel standard which can be switched to form a network. This architecture is currently the only way to get true gigabits per second delivered to a machine.

There are several ways to hedge the high performance network bet. Other than FDDI, the emerging standards mentioned above are still under develop ment and may not survive. Therefore, many will want to wait until the network war settles down.

Another recommendation concerns the physical layer. In order to support a gigabit infrastructure, a proper cable plant needs to be installed and it can be installed now ! The next sections of this paper discuss physical layer options and recommendations.

5 Physical Layer Concentration

The physical layer is one of the most important key elements in the provision of gigabit architectures. Planning and implementation start by considering the answer to one major question . Who needs to be connected and what are the distance distributions that must be considered? This one question raises other issues such as , are there external networks to be considered, is the physical plant transparent enough to provide connectivity for today's technologies and is there capacity for expansion of new applications and technology hardware? A physical plant could be as simple as an internal office connection between three personal computer systems located within 100 feet of each other or as complex as many buildings with many users and varying applications separated by reasonably good distances. In many situations, the enterprise can structurally look like a small community which calls for a robust physical plant delivering a variety of services and appropriate bandwidths.

A physical layer plant consists of two main components; the Outside Plant (OSP) and the Inside Plant (ISP). The structure of the physical layer plant consists of many elements. The main elements which comprise a physical plant are; cables which connect the main buildings of the enterprise typically called the buried jelly filled (BJF) distribution, connectors (terminations) for the OSP cables, the Inside Plant (ISP) wiring, ISP interconnect cabinets, ISP cabling connectors and station jacks (terminations), plant pathways (OSP and ISP) and records, documentation and administration. Local codes, ordinances, Federal regulations and codes , and industry standards will help to shape the compliance and overall installation of the physical plant elements.

6 Outside Physical Plant

The OSP physical plant cabling system must provide connectivity to each enterprise building that will require network connectivity for gigabit speeds. Copper cabling systems become limited in installation applications to compliment gigabit transmission hardware. High bit rate OSP cables do not typically exist. Transmissions on average 24 gauge buried jelly filled (BJF) multiple pair copper cable can only provide reasonable quality by utilizing conservative designs using horizontal displacement of balanced transmit and receive pairs. Repeaters are also placed at short intervals. Buried or manhole repeaters also require the provision of span power. Copper OSP transmission media becomes limited at speeds just above T1 rates (1.544 Mbps). Gigabit transmissions start at 1000 Mbps.

The next likely candidate for transmission becomes coaxial cable. Coaxial cable has greater bandwidth carrying capabilities. It has a full 360 degree rotational shield that minimizes the effects of electromagnetic intrusions. Coaxial cable, however, does not provide a good cost/performance ratio. Additionally, emerging standards are increasingly less likely to support uses of coaxial transport backbones. Comparisons to industry high bandwidth service providers like the cable television industry are indicating that coaxial cable plants are being replaced with single mode fiber optic cables.

Fiber optic cables have been on the increase in all enterprise OSP backbones. Fiber optic cables are not effected by comparative problems found in copper cable plants. Fiber optic spans are typically designed as continuous cable runs. Where necessary, fiber optic cables can be spliced together with additional cable or to be used as a repair medium for cable breaks due to backhoes and other unplanned activities that can damage cables such as rodents. Fiber optic cables have greater information carrying capacities and can run greater distances before needing to be repeated. Fiber optic cable is lighter and physically smaller in size as compared to copper cable. Fiber optic cables are immune from the effects of electromagnetic and electrostatic interference. The installation, termination and maintenance of fiber optic cables however, requires increased learning/training activities, parts inventories, tools and test equipment.

7 Inside Physical Plant

ISP cabling systems emanate from two camps; the copper world (shielded and unshielded) and the fiber optic world (multimode and singlemode). High bit rate (up to 100Mbps) copper distribution cables can contain up to 4 pairs in a shielded or unshielded sheath. Fiber optic cables are of two types; constructed as multiple strand multimode or singlemode fibers or combinations of both known as hybrid cables.

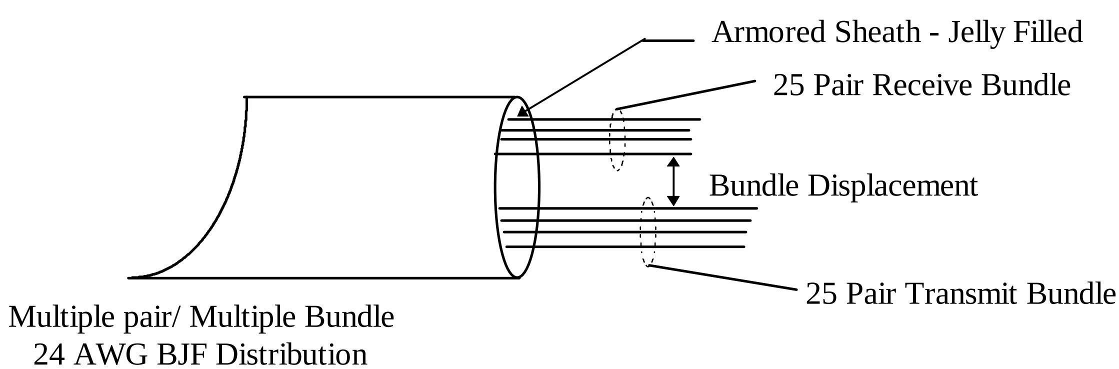

Split Pair Bundle Transmission

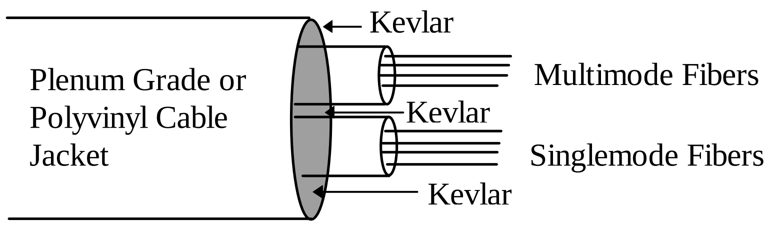

Hybrid Fiber Optic Cable

Current copper high bit rate transmission strategies endorse the use of unshielded twisted pairs. High bit rate technologies specify the use of 4 pairs of copper wire under a common sheath. These types of cables are known as unshielded twisted pair (UTP) cables and are becoming commonly used in Category 5 (also known as Level 5), 100 MegaBit per second (Mbps) transmission applications. The category system is a way of rating cables by their ability to be compliant to a set of criteria for transmission.

One key element to the design of Category 5 cables are the twist technologies that are used. Each pair of wires has a unique amount of twists per foot of cable. The twisting of the pairs allows for a 180 degree phase cancellation of any unwanted or induced noises. This is an important factor to consider when multiple pairs within the cable may contain different signaling or the cable may reside close to another signal carrying cable. The ratio between signal and noise (the signal to noise ratio) should remain as large as possible to maximize the quality and undisturbance of the original signal and to minimize any noise. This is a formidable engineering task when there is no shield around wire pairs. Theoretically if one were to keep increasing the twists per foot on a single pair of wires, at some point the cable would turn into a coaxial cable.

UTP cables are also susceptible to crosstalk. Crosstalk is simply the bleeding over or inducing of one cable's signal to another pair of wires in the same or another cable. Crosstalk can distort or disguise the original information making it unintelligible at the distant end. Near End crosstalk (NEXT) has become an important factor in the selection of high bit rate cables. Standards committees have incorporated acceptable ranges and criteria for NEXT measurements. There are now several test instruments emerging on the market which can test cables for losses, crosstalk and Category 5 performance.

Another important factor in the construction of Category 5 cabling is the control of attenuation. Attenuation can be controlled by maintaining copper wire diameter and insulation thickness around the copper wire. Characteristic impedance is also of related importance. Maintaining a proper characteristic impedance between the source generator and the destination receiver will ensure maximum power transfer and minimizes the possibility for distortion. Points along the cable path that are inconsistent in impedance produce reflections which equates to signal loss. Keeping the signal at a proper transfer level will ensure that the signal is reproducible when inputted at the far end amplifier. As the amplifier increases the amplitude of all the signals received, it is important to keep noise at a minimum otherwise the noise becomes amplified and the o utput signal becomes unintelligible. The key to using these high bit rate cables successfully is to follow very specific installation guidelines which allow the cable to function with maximum cable performance.

The strength to building future networks with unshielded twisted pair cable is that it is a familiar technology and there is a wide base of supporting manufacturers. Installations personnel do not find it difficult to strip the cable back and terminate the cable pair ends. Care must be taken, when terminating cable pairs, to maintain pair twists to within one half inch of the point of termination. Failure to maintain pair twists can jeopardize system performance. Plenum rated cables (CMP) are more difficult to terminate, but installation personnel can acquire the proper touch for stripping the Teflon jacketing on these cables. UTP polyvinyl jacketing (CMR rated) is flexible but yet strong enough to protect the cable pairs.

UTP cable performs well at various transmission bit rates and frequencies up to 100 MHz (some now claim 350 MHz) within a 90 meter distribution limit. Additionally, certain manufacturers are now specifying that their product s support cable transmission bit rates of 155 Mbps (e.g. Asynchronous Transfer Mode - ATM). The actual transmission bit rate is dependent upon the encoding algorithm used.

Category 5 cable pairs are color coded with standard telephone wire colors. There is a wide variety of termination hardware to choose from. Cable can be easily removed from pathways when necessary. There are many manufacturers that produce these types of cables. Manufacturers' natural competition drives the market which keeps increasing the overall product quality, performance and cost.

There is a wide variety of manufacturers that produce termination hardware for copper wire cabling. Category 5 products include a wide variety of station jacks with several options for termination. Two of the most popular termination schemes are the Electronics Industry Association (EIA) and Telecommunications Industry Association (TIA) 568A and 568B. These standards spell out specifications for Category 5 wiring including which wire pairs terminate on the particular pin numbers of a n 8 pin modular jack. This is sometimes known as the sequence.

The 8 pin modular jack is also known as an RJ45, although RJ45 is technically incorrect. The name RJ45 or RJ45S was used years ago for an 8 pin jack with two wires terminating on it (ring on pin 4 and tip on pin 5) and a programming resistor between pins 7 and 8. The name some how stuck with this type of jack and is used (including manufacturers) to reference the 8 pin modular jack. The EIA/TIA 568 standard specifies the jack as a n "8 pin modular".

There are several limitations to high bit rate copper cabling systems. Companies have strived to produce the best cable. In producing this type of cable, companies have had to adjust their design and produce revision level cables to accommodate increases in performance design as well as changes in Standards. For example, a manufacturers first product release was found not to install as smoothly because the cable had a tendency to kink due to tough material in the jacketing. Later revisions cured this jacketing problem. Additionally, it is not believed at the date of this writing that enterprise Category 5 cabling networks exist. Manufacturers have extensively tested their products under strict laboratory conditions with highly sophisticated and extremely expensive test equipment. Manufacturers have also ensured that their test installations have been meticulously installed and tweaked to perfection. Today, there are a few manufacturers that are commercially marketing Category 5 test equipment. These testers may range in price from $2,000 to $4,000. A disadvantage is that test results may not always be consistent. Also the time to test a cable with a commercial tester, in only one direction, may take five or more minutes. This may be a concern when the enterprise technical staff is performing tests on an installation (in both directions) with several hundred jacks and users want their system available to them. Laboratory test gear costs approximately $35,000. Are there any important differences between the commercially available test gear and the laboratory test gear?

Enterprises who have bought early into certain wiring technologies and equipment may be limited to developing changes in transmission requirements and limitations or obsolescence of early production product. It is not necessarily a case of waiting for technology to develop, but rather to perform discovery in the market technology and think through all aspects before implementing partial decisions. A quality system is a combination of a reliable and expandable configuration, reliable components, accurate installation, good labeling, reliable test equipment, a trained staff and excellent documentation.

The Federal Communications Commission has also set restrictions on Category 5 cable industry specifications. Main concentrations of energy must remain below a spectral frequency of 30 MHz. This principle is enacted to keep the wire from acting as a transmission antenna in which it can interfere with other electrical products and frequencies. This also ensures noninterference with the public safety frequency band. It therefore becomes the task of transmission equipment manufacturers to develop complex encoding algorithms to increase throughput. One example of a four level coding algorithm is 2B1Q. One of the newer more complex multilevel coding algorithms is CAP32.

Fiber optic cable is an excellent choice for enterprise backbones. Whether the backbone is OSP or ISP, fiber optic cable eliminates distance/distribution problems which are a typical nemesis of copper cable systems. Fiber optic cables are also starting to accompany copper station cables in horizontal distributions to the desktop. Fiber optic cables are free from electromagnetic intrusions, are quite stable and present a more secure means for priority transmissions. When matched with the proper mode of operation, fiber optic cables have exceptional ability to accommodate transmission bandwidths. Fiber optic cable in this paper refers directly to glass (silica) optical fiber strands. Plastic clad silica (PCS) based optical fiber has low bandwidth capabilities and high losses and is not herein considered.

There are basically two types of fiber optic cable, multimode and singlemode. This paper considers the merits of graded index multimode and step index singlemode fiber cable plants. Both multimode and singlemode fiber optic cable can be manufactured as loose tube or tight buffered types. These types of cable constructions have to do with how the fibers are packed within the sheath. Loose tube cable contains plastic material tubes which contain a specified amount of fiber strands per tube. The tubes are surrounded by Kevlar strength members and covered by the sheath. The tube always leaves plenty of room for the fibers to move about freely. This allows for expansion /contraction of fiber in the cable d ue to thermal expansion and is typically used in OSP installations. Loose tube cables also contain a buffering material on each fiber acting as an extra coating over the cladding. A typical buffer for OSP fibers is 250 um. Tight buffered cable does not contain any tubes. Fiber strands are also covered by a buffer, typically 900 um, and packed with Kevlar under the cable sheath. The thicker buffer adds stability to the cable especially when some ISP cables do not contain a metal encasement as part of the sheath structure. Kevlar gives added protection to the cable by keeping the fibers resistant to impact and flexing of the cable during installation and sharp turns. Tight buffered cable is used in vertical and horizontal ISP installations.

Fiber optic cable consists mainly of two parts; the core and the cladding. The cladding surrounds the light carrying portion of the fiber (core) and helps to concentrate light waves traveling in different modal paths. One of the most popular core sizes for multimode cable is 62.5 um. Singlemode cable has popular core sizes between 8.3-9 um. Common claddings are 125 um used in both multimode and singlemode cable.

A characteristic of multimode fiber is that pulses entering the cable become wider and rounded by the time they exit the cable. This effect is known as pulse dispersion. There are two type of dispersion; modal and chromatic. Dispersion places a limit on the amount of cable bandwidth. Pulses entering the cable must be narrow enough that they do not overlap by the time that they exit the cable. If pulses do overlap, the result is indistinguishable data at the receiver end. Singlemode, however, has increased bandwidth compared to multimode because the effects of dispersion are minimized. Multimode cable, however, is widely used and gives good cost/performance characteristics.

Hybrid fiber optic cables are also being produced which contain both multimode and singlemode fibers. Using a hybrid cable can decrease cable material costs because less cabling material is required as well as man hours necessary for installation.

Cable manufacturers are now marketing cables that contain a mixture of fiber optic cables and/or copper. Typical OSP hybrid cables are constructed with multimode and singlemode fibers. In the ISP, hybrids are also constructed from multimode and singlemode fibers. There are also cables that contain multimode fiber, singlemode fiber and Category 5 copper. These cables are actually three sheaths laid out side by side and bonded together. These three compartment type cables are sometimes called, “triamese” cables. Manufacturers are trying to construct a regular rounded form cable. They are also trying to add a second Category 5, four pair cable. The advantage of hybrid cables is that all fibers are concentrated in one cable and therefore the associated cost to install the cable is less. Sometimes unions charge rates that are based on the actual physical cable count. The disadvantage to the three compartment triamese design is that the cable is more of a rectangular shape instead of being rounded and it becomes harder to install in conduit. Since the outer compartments of the cable bend easily, the cable has more surface area to get caught up in the installation path. The cable gets hung up easier on other cables in cable trays, especially if the tray is already packed with other types of communications cables.

There is a wide variety of transmission equipment for multimode and singlemode fiber. Transmission equipment with multimode interfaces is reasonably priced , especially if a light emitting diode (LED) driver/receiver is selected. If a fiber optic cable plant becomes overloaded with applications, there are other transmission options to take better advantage of the existing cable before more cable must be added to the cable plant. Wave division multiplexing (WDM) equipment can be used to allow a single pair of fibers to transmit/receive two independent applications, each residing at a fixed transmission wavelength. WDM equipment can be expensive and the cost of installing more fiber can seem more acceptable.

Singlemode fiber is a more complex media to work with. Installation personnel need increased training to recognize the differences and challenges presented by singlemode. Singlemode cable has slightly increased in cost, takes more attention for termination, has greater bandwidth and is driven by higher cost transmission equipment. Installation personnel must be more careful not to create macrobends by crushing the cable with tie wraps when installing singlemode cable. Proper test equipment is a key to maintaining a fiber optic plant. Singlemode fiber is the key to building future gigabit networks which will seed new application horizons. Singlemode fiber eliminates network throughput speed bumps experienced in other technologies.

8 Fiber Optic Connectors

There is a wide variety of companies making terminating connectors for fiber optic cable. The most popular connector types are; Bionic, SMA, ST, FC, FSD and SC. Fiber optic connectors can be the least common denominator in the performance of the cable plant. It is important to pick connectors that are easy to work with as well as presenting least attenuation. When connecting one fiber optic cable strand to another, cross connect frames are typically used. These frames consist of cable strands terminated with fiber optic connectors which are interconnected by fiber optic couplers.

Fiber optic couplers allow terminated cables to be connected together to other terminated cables to make network circuit paths between different buildings or locations. New systems have entered the market which are called connectorless systems in that they do not use any of the above termination connector types. These connectorless systems insert fiber into splice-like modules to make cross connections. The modules can be reused a number of times and when necessary they can be rebuilt to restore original ability. The only places where connectors are necessary are at the originating and destination ends of the plant.

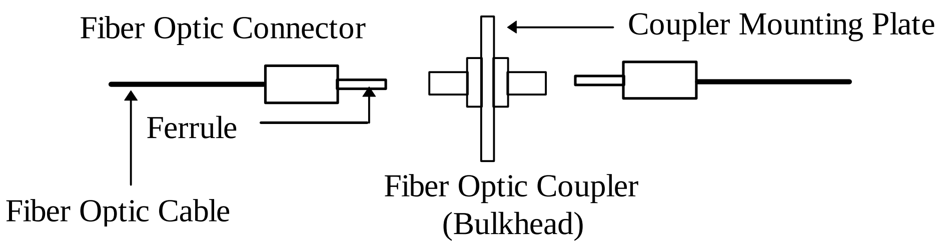

Traditional Interconnection of Fiber Optic Cables

Couplers, (sometimes called bulkheads) are designed to allow a fiber optic cable strand to be connected to each side of the coupler thus performing an interconnection between the two fiber strands. When this is done light leaves the glass , passes through a small air gap and enters the fiber optic connector on the other side of the coupler. The changing from one index to another (glass to air) causes a small percentage of light to be reflected back into the emanating fiber strand. This causes a transmission loss. Many connectors or cross connects in a fiber path can inhibit the system's overall ability. Connectorless systems along with mechanical splices try to circumvent changes in the index of refraction by using an index matching gel between similar indexed materials. The advantages to a connectorless system is lower attenuation, reduction in reflections and increased system performance.

9 Argonne's Gigabit Architecture

Argonne National Laboratory has recently deployed a new design for its Laboratory-wide network infrastructure. This infrastructure has enabled Argonne to implement gigabit architectures as well as to accommodate all current computer communication needs.

The Argonne outside plant consists of 24 major buildings connected in a dual starred topology. Fiber optic cables from two main hub buildings connect to other campus buildings. One hub is located at the North region of the campus and one hub is located in the South region of the campus. Installed fiber optic cables are of a hybrid cable design that consists of 24 graded index multimode fibers and 12 step index singlemode fibers. The cable is a loose tube type design with a jelly filled water retardant. OSP cables have an armored type jacketing. Armored sheaths under the jacketing provide cable stability as well as rodent proof protection. Yes, rodents will try to eat up the cable plant especially if the cable is direct buried.

The main link between hubs consists of two cables of the same 24/12 hybrid mixture. Most all of the OSP cables are installed in 4 inch schedule 40 PVC pipe except for a few direct burials. Direct burials are buried at a minimum of thirty inches. Generally a copper cable is buried with the fiber so that the cable can be traced with a cable locator.

All cables terminate on high density connector cross connect panels. The cross connect panels are of two types; wall mount and rack mount. The panels offer a unique opportunity compared to other termination hardware because a large number of terminations can be housed in a very small physical area. OSP cables terminate on one side of the high density connector. The OSP cable strands are 250 um buffered. All cross connects are performed with 900 um buffered cable. The cross connect cable is a transparent white colored cross connect. This color allows for a fiber optic light source to be used to service the cable plant and detect broken cross connects. The design of the high density connectors accommodate a mixture of buffer sizes. In Argonne's case it is 250 um for outside plant and 900 um for cross connects and horizontal distribution.

One major advantage of high density connectors is that attenuation is less than conventional fiber optic connectors. Connectors contain an index matching gel to circumvent any air gaps created between fibers during installation. This may be a critical advantage when there are many hops or cross connects in a plant's fiber optic circuits especially in achievement of gigabit speeds across singlemode cable.

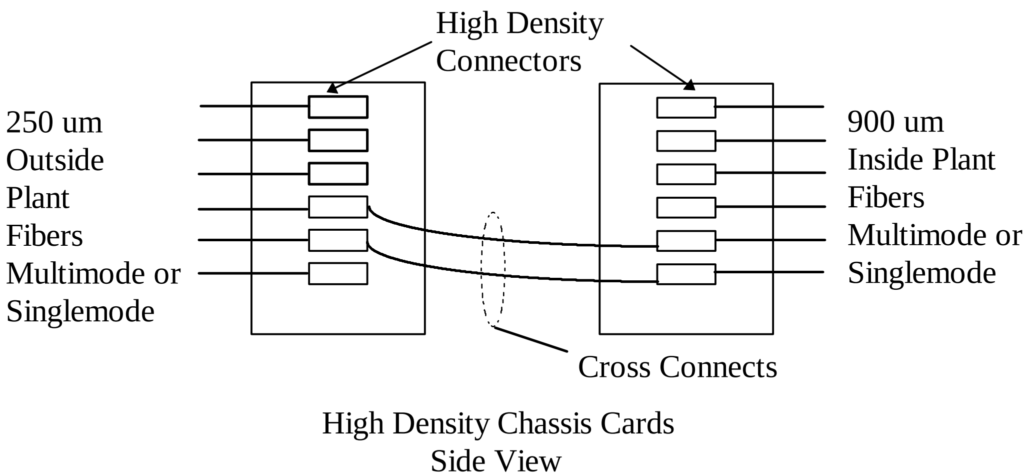

Fiber Plant Termination and Cross Connect

ISP fiber optic riser cables route to other cross connect panels or directly to user application areas. Argonne's gigabit architecture has been given direction from some typical applications infrastructure. Typical ISP horizontal distributions for FDDI and bridged Ethernet connectivity consist of six terminated multimode strands in a single cable. Cables are terminated with fiber termination modules (FTM) which fasten into a common housing. FTMs are preterminated with an ST type connector linked by a fiber strand to a mechanical splice. The other side of the mechanical splice is attached to the fiber optic horizontal distribution cable. The FTMs are excellent for field terminations where a minimum of tools and equipment are necessary. The packaging is also of importance as users of these higher speed systems prefer small aesthetically pleasing communications outlets. Successful implementation of these services has provided direction to Argonne's comprehensive gigabit architecture.

The recommended Gigabit configuration consists of 2 multimode fibers with a 62.5 um core and 125 um cladding, 2 singlemode fibers with a core of 8-9 um and a cladding of 125 um and 1 or 2 optional copper Category 5 or higher, 4 pair UTP cables. ST type connectors are used on the multimode cable and SC type connectors are used on the singlemode cable. SC connectors allow for higher performance where singlemode cable is more difficult to work with. Using two connectors allows personnel to identify any termination in the plant and correctly determine whether the cable is singlemode or multimode cable.

Future revisions in cable construction technologies may allow for a hybrid ISP cable containing 2 Category 5 cables along with multimode and singlemode fibers in a concise rounded cable jacket.

The following diagram illustrates the last point of distribution to the workstation/office area. A regional closet with good pathway access can be utilized to hold Category 5 wiring panels and also fiber optic panels.

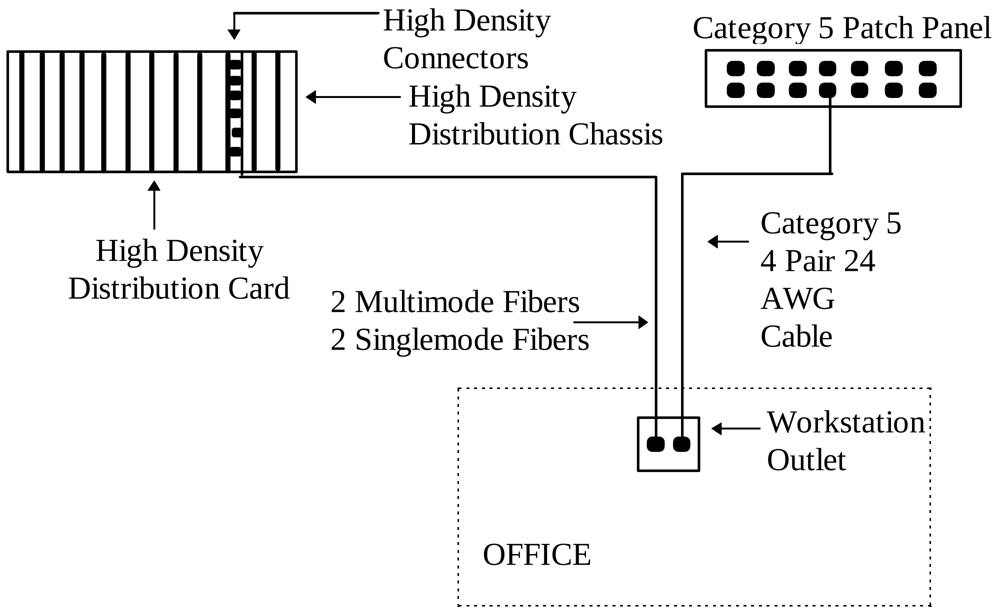

Gigabit Distribution to Office Environment

10 Conclusion

This paper briefly discusses the various considerations to deploying gigabits to the desktop. In the view of the authors, the most important enabler of this goal is a solid high speed infrastructure. Fiber optic media is the only one capable of supporting this infrastructure. Cable and hardware manufacturers need to demystify fiber optic technology before this can happen on a large scale. A goal of demystifing fiber optic technology and associated products can be achieved through easy and inexpensive install ation solutions as compared to high performance copper systems. This must absolutely be achieved before the general population begins to consider fiber optic options. Until that time, the deployment of gigabits to the desktop will remain limited.